arXiv:0709.0670v1 [cs.DS] 5 Sep 2007

Using Data Compressors to Construct Rank Tests

Daniil Ryabko, Jürgen Schmidhuber
IDSIA, Galleria 2, CH-6928 Manno-Lugano, Switzerland
{daniil, juergen}@idsia.ch

Abstract

Nonparametric rank tests for homogeneity and component independence are proposed, which are based on data compressors. For homogeneity testing the idea is to compress the binary string obtained by ordering the two joint samples and writing 0 if the element is from the first sample and 1 if it is from the second sample and breaking ties by randomization (extension to the case of multiple samples is straightforward). H_0 should be rejected if the string is compressed (to a certain degree) and accepted otherwise. We show that such a test obtained from an ideal data compressor is valid against all alternatives.

Component independence is reduced to homogeneity testing by constructing two samples, one of which is the first half of the original and the other is the second half with one of the components randomly permuted.

1 Introduction

We consider two classical problems of mathematical statistics. The first one is homogeneity testing: two (or more; see below) samples $X_1, \dots X_n$ and $Y_1, \dots, Y_n$ with elements in $\mathbb{R}$ are given. It is assumed that the elements are drawn independently and within samples the distribution is the same. We want to test the hypothesis H_0 that X_i and Y_i are distributed according to the same distribution versus H_1 that the distributions generating the samples

are different. This is called homogeneity testing. Absolutely no assumptions are made on the distributions.

The second one is component independence: a sample $Z_1, \ldots, Z_n$ is given, generated i.i.d. according to some distribution F_Z. Each element Z_i consists of two (or more) components Z_i^1 and Z_i^2. We wish to test whether the components are independent of each other. That is, H_0 is that the marginal distributions are independent whereas H_1 is that there is some dependency. Again, no assumption is made on the distribution F_Z.

Both problems are well-known problems of nonparametric mathematical statistics. For example, a classical test for homogeneity is Kolmogorov-Smirnov test (which assumes, however, that the distributions generating the samples are continuous). There are many other nonparametric tests; some of the tests use *ranks* of elements within the joint sample, instead of using the actual samples. Such is, for example, Wilcoxon's test, see [3] for an overview (which also makes some additional assumptions on the distribution).

In this work we present simple nonparametric (distribution-free) rank tests for homogeneity and component independence based on data compressors.

The idea to use real-life data compressors for testing classical statistical hypotheses, such as homogeneity, component independence and some others, was suggested in [7, 8]. In these works statistical tests based on data compressors are constructed which fall into the classical framework of nonparametric mathematical statistics, in particular, the Type I error is fixed while Type II error goes to 0 under a wide range of alternatives. The hypotheses considered there mostly concern data samples drawn from *discrete* (e.g. finite) spaces. Some tests for continuous spaces are also proposed based on partitioning. Here we extend this approach to *rank* tests, allowing testing homogeneity and component independence without the need of partitioning the sample spaces and making them finite. The idea of using data compressors for tasks other than actual data compression was suggested in [1, 2, 4], where data compressors are applied to such tasks as classification and clustering. These works were largely inspired by Kolmogorov complexity, which is also an important tool for the present work.

An "ideal" data compressor is the one that compresses its input up to its Kolmogorov complexity. This is intuitively obvious since, informally, Kolmogorov complexity of a string is the length of the shortest program that outputs this string. Such data compressors do not exist; in particular, Kolmogorov complexity itself is incomputable. Real data compressors, however,

can be considered as approximations of ideal ones.

In this work we provide a simple empirical procedure for testing homogeneity and component independence with data compressors; we show that for an ideal data compressor this procedure provides a statistical test which is valid against all alternatives (Type II error goes to zero); while Type I error is guaranteed to be below a pre-defined level (so-called significance level) for all data compressors, not only for ideal ones. It should also be noted that the theoretical assumption underlying data compressors used in real life is that the data to compress is *stationary*. Thus the tests designed in [7, 8] are provably valid against any stationary and ergodic alternative, while these tests are based on real data compressors, not only on ideal ones. In our case, the alternative arising in rank test under H_1 is not stationary. Thus we prove theorems only about ideal data compressors, and real data compressors can be used heuristically. However, it can be conjectured that the same results can be proven for some particular real-life data compressors, for example for those which are based on the measure R from [6] or on the LZ algorithm [10].

2 Homogeneity testing

Homogeneity testing is the following task. Let there be given two samples $X = \{X_1, \dots X_m\}$ and $Y = \{Y_1, \dots, Y_k\}$ (the case of more than two samples will also be considered). X_i are drawn independently according to some probability distribution F_X on $\mathbb{R}^d$ ($d \in \mathbb{N}$) and Y_i are drawn independently from each other and from X_i according to some distribution F_Y on $\mathbb{R}^d$. The goal is to test whether $F_X = F_Y$. No assumption is made on the distributions F_X and F_Y; we only assume that X_i and Y_i are drawn independently within the samples and jointly. So, we wish to test the hypothesis $H_0 = \{(F_X, F_Y) : F_X = F_Y\}$ against $H_1 = \{(F_X, F_Y) : F_X \neq F_Y\}$.

A *code* φ is a function $\varphi : B^* \to B^*$ from the set of all finite words over binary alphabet $B = \{0, 1\}$ to itself, such that φ is an injection (that is, $a \neq b$ implies $\varphi(a) \neq \varphi(b)$ for $a, b \in B^*$). A trivial example of a code is the identity $\varphi_{id}(a) = a$. Less trivial examples that we have in mind are data compressors, such as `zip, rar, arj,` or others, which take a word and output a "compressed" version of it (which in fact is often longer than the original) from which the original input can always be recovered. We will construct (reasonable) tests for homogeneity from (good) data compressors.

First let us assume that $\mathbf{d} = \mathbf{1}$ (that is, $X_i, Y_i \in \mathbb{R}$). Let $Z_1 \leq Z_2 \leq \dots \leq$

Z_{m+k} denote the joint sample constructed by ordering jointly two samples X and Y. Construct the word $A = A_1 \ldots, A_{m+k}$ as follows: for each i $A_i = 0$ if Z_i is taken from the sample X ($Z_i \in X$) and $A_i = 1$ if Z_i is from the sample Y ($Z_i \in Y$) where ties are broken by randomization: if $Z_j = Z_{j+1} = \ldots Z_{j'}$ and there are m' elements of the sample X which are equal to Z_j and k' elements of the sample Y which are equal to Z_j then the word $A_j \ldots A_{j'}$ is chosen randomly from all $\frac{(m'+k')!}{m'!k'!}$ binary words which have m' zeros and k' ones, assigning equal probabilities to all words.

Now consider the case $\mathbf{d > 1}$, that is, the elements of the samples X and Y are from $\mathbb{R}^d$, $d > 1$. Construct samples $\bar{X} = \bar{X}_1, \ldots, \bar{X}_m$ and $\bar{Y} = \bar{Y}_1, \ldots, \bar{Y}_m$ as follows: $\bar{X}_t := x_t^{11}, x_t^{21}, \ldots, x_t^{d1}, x_t^{12}, x_t^{22}, \ldots, x_t^{d2}, \ldots$ where x_t^{ij} is the jth element in the binary expansion of the ith component of X_t (in case the expansion is ambiguous always take the one with more zeros), and analogously for Y. Denote the described function which converts X to $\bar{X}$ by τ. Construct the string A applying the (single–dimensional) procedure described above to the samples $\bar{X}$ and $\bar{Y}$.

Let $|K|$ denote the length of a string K.

Definition 1 (Homogeneity test G_φ)**.** *For any code φ the test for homogeneity G_φ is constructed as follows. It rejects the hypothesis H_0 (outputs* reject*) at the level of significance α if*

$$|\varphi(A)| \leq \log \alpha N \tag{1}$$

where $N := \frac{(m+k)!}{m!k!}$ and log *is base 2, and accepts H_0 (outputs* accept*) otherwise.*

Definition 2 (More than two samples)**.** *In case we are given r samples where $r \geq 2$ and wish to test H_0 that they all are generated according to the same distribution versus at least two distributions are different, the test is the same, except for that the string A is not binary but from r-element alphabet and in the test above instead of N take*

$$N' := \frac{(\sum_{i=1}^r m_i)!}{\prod_{i=1}^r m_i},$$

where m_i are the sizes of the samples.

The intuition is as follows. Observe that if the distributions F_X and F_Y are equal (that is, H_0 is true), then the string A is just a random binary string

with m zeros and k ones; all such strings have equal probabilities under H_0. Thus a good data compressor should be able to compress it to about $\log N$ bits, but no code can compress many such strings to less than $\log N - t$ bits ($t > 0$), since there are N such strings and only $2^{-t}N$ binary strings of length $\log N - t$.

Proposition 1 (Type I error)**.** *Let $d = 1$. For any code φ and any $\alpha \in [0, 1]$ the Type I error of the test G_φ with level of significance α is not greater than α:*

$$P\{X, Y : G_\varphi(X, Y) = reject\} \leq \alpha \tag{2}$$

for all $P = (F_X, F_Y) \in H_0$.

Remark 1. *The proposition still holds if H_0 is rejected when*

$$|\varphi(A)| \leq (k+m)h\left(\frac{k}{k+m}\right) + \log\alpha - \log(k+m), \tag{3}$$

where $h(t)$ is the entropy

$$h(t) := -t\log t - (1-t)\log(1-t). \tag{4}$$

In case of r samples (3) takes the form

$$|\varphi(A)| \leq nh\log r + \log\alpha - \log n \tag{5}$$

with $n = \sum_{i=1}^{r} m_i$ and $h = -\sum_{i=1}^{r} \frac{m_i}{n}\log\frac{m_i}{n}$.

Proof. As it was noted, under H_0 for every string $a \in B^{k+m}$ such that a consists of m zeros and k ones $P(A = a) = 1/N$ (that is, all such strings are equiprobable). Since there are only αN binary strings of length $\log \alpha N$ and φ is an injective function, that is each codeword is assigned to at most one word, we get $P\{X, Y : |\varphi(A)| \leq \log \alpha N\} \leq \frac{1}{N}N\alpha = \alpha$ which together with the definition of G_φ implies (2).

The statement of the Remark can be derived from Stirling's expansion for N and N'. □

Remark 2. *The term $-\log(k+m)$ in (3) is due to the fact that there are only $\frac{(m+k)!}{m!k!}$ strings with m zeros and k ones (among 2^{k+m} all binary strings of this length). So the code φ can specifically assign shorter codewords to these strings. As real data compressors are not designed to favour strings of this particular ratio of zeros and ones, in practice it is recommended to omit the term $-\log(k+m)$ in (3). The same concerns the term $-\log n$ in (5).*

Obviously, for some codes the test is useless (for example if φ is the identity mapping) and Proposition 1 is only useful when the Type II error goes to zero. Next we will define "ideal" codes (the codes that compress a word up to its Kolmogorov complexity) and show that for them indeed the probability of *accept* goes to zero under any distribution in H_1.

Informally, Kolmogorov complexity of a string A is the length of the shortest program that outputs A (on the empty input). Clearly, the best, "ideal", data compressor can compress any string A up to its Kolmogorov complexity, and not more (except may be for a constant). Next we present a definition of Kolmogorov complexity; for fine details see [9, 5]. The complexity of a string $A \in B^*$ with respect to a Turing machine ζ is defined as

$$C_\zeta(A) = \min_p \{l(p) : \zeta(p) = A\},$$

where p ranges over all binary strings (interpreted as programs for ζ; minimum over empty set is defined as ∞). There exists a Turing machine ζ such that $C_\zeta(A) \leq C_{\zeta'}(A) + c_{\zeta'}$ for any A and any Turing machine ζ' (the constant $c_{\zeta'}$ depends on ζ' but not on A). Fix any such ζ and define *Kolmogorov complexity* of a string $A \in \{0,1\}^\infty$ as $C(A) := C_\zeta(A)$. Clearly, $C(A) \leq |A| + b$ for any A and for some b depending only on ζ.

Definition 3 (ideal codes)**.** *Call a code φ* ideal *if some constant c the equality $|\varphi(A)| \leq C(A) + c$ holds for any binary string A.*

Clearly such codes exist.

Proposition 2 (Type II error: universal validity)**.** *For any ideal code φ Type II error of the test G_φ with any fixed significance level $\alpha > 0$ goes to zero $P\{X, Y : G_\varphi(X, Y) = accept\} \to 0$ for any P in H_1 if $k, m \to \infty$ in such a way that $0 < a < \frac{k}{m} < b < 1$ for some a, b.*

Proof. First observe that the function τ that converts d-dimensional samples X and Y to single-dimensional samples $\bar{X}$ and $\bar{Y}$ has the following properties: if X and Y are distributed according to different distributions then X and Y are also distributed according to different distributions. Indeed, τ is one to one, and transforms cylinder sets (sets of the form $\{x \in \mathbb{R}^d : x^{i_1 j_1} = b_1, \ldots, x^{i_t j_t} = b_t; b_l \in \{0,1\}, t, i_l, j_l \in \mathbb{N}(1 \leq l \leq t)\}$) to cylinder sets. So together with F_X (F_Y) it defines some distribution $F_{\bar{X}}$ ($F_{\bar{Y}}$) on $\mathbb{R}$. If distributions F_X and F_Y are different then they are different on some cylinder set

T, but then $F_{\bar{X}}(\tau(T)) \neq F_{\bar{Y}}(\tau(T))$. Thus further in the proof we will assume that $d = 1$.

We have to show that Kolmogorov complexity $C(A) = |\varphi(A)|$ of the string A is less than $\log \alpha N \geq (k+m)h\left(\frac{k}{k+m}\right) + \log \alpha - \log(k+m)$ for any fixed α from some k, m on. To show this, we have to find a sufficiently short description $s(A)$ of the string A; then the Kolmogorov complexity $|\varphi(A)|$ is not greater than $|s(A)| + c$ where c is a constant.

If H_1 is true then $F_X \neq F_Y$ and so there exist some interval $T = (-\infty, t]$ and some $\delta > 0$ such that $|F_X(T) - F_Y(T)| > 2\delta$. Then we will have

$$\left| \frac{\#\{x \in X \cap T\}}{m} - \frac{\#\{y \in Y \cap T\}}{k} \right| > \delta \tag{6}$$

from some k, m on with probability 1.

Let A' be the starting part of A that consists of all elements that belong to T and let $m' := \#\{x \in X \cap T\}$ and $k' := \#\{y \in Y \cap T\}$. A description of A' can be constructed as the index of A' in the set (ordered, say, lexicographically) of all binary strings of length $m'+k'$ that have exactly m' zeros and k' ones plus the description of m' and k'. Thus the length of such a description is bounded by $\log \frac{(k'+m')!}{k'!m'!} \leq (m'+k')h(\frac{k'}{m'+k'})$ plus $\log k' + \log m' + const$ (the inequality follows from $n! \leq n^n$ for all n). Let $\bar{A}$ denote the remaining part of A (that is, what goes after A'). The length of the description of $\bar{A}$ is bounded by $(\bar{m}+\bar{k})h(\frac{\bar{k}}{\bar{m}+\bar{k}}) + \log \bar{k} + \log \bar{m} + const$ where $\bar{m} = m - m'$ and $\bar{k} = k - k'$. Since h is concave and $\frac{k}{m+k}$ is between $\frac{k'}{m'+k'}$ and $\frac{\bar{k}}{\bar{m}+\bar{k}}$, from Jensen's inequality we obtain

$$h\left(\frac{k}{m+k}\right) - \left(\frac{m'+k'}{m+k} h\left(\frac{k'}{m'+k'}\right) + \frac{\bar{m}+\bar{k}}{m+k} h\left(\frac{\bar{k}}{\bar{m}+\bar{k}}\right)\right) > 0.$$

Denote this difference by $\gamma(k, m, k', m')$. Let $\gamma = \inf \gamma(k, m, k', m')$ where the infimum is taken over all pairs k, m that satisfy the condition of the proposition $0 < a < \frac{k}{m} < b < 1$ and k', m' that satisfy (6). It follows that $\inf |\frac{k'}{m'} - \frac{k}{m}| > 0$ and $\inf |\frac{\bar{k}}{\bar{m}} - \frac{k}{m}| > 0$. Thus, γ is positive and depends only on a, b and δ. To uniquely describe A we need the description of A' and $\bar{A}$ and also k and m; these have to be encoded in a self-delimiting way; the length of such a description $s(A)$ is bounded by the lengths of description of A', $\bar{A}$

plus $\log(k+m)$ and some constant. Thus

$$(k+m)h\left(\frac{k}{k+m}\right)+\log\alpha-\log(k+m)-|\varphi(A)|\geq$$
$$(k+m)h\left(\frac{k}{k+m}\right)+\log\alpha-2\log(k+m)$$
$$-\left(\frac{m'+k'}{m+k}h\left(\frac{k'}{m'+k'}\right)+\frac{\bar{m}+\bar{k}}{m+k}h\left(\frac{\bar{k}}{\bar{m}+\bar{k}}\right)\right)-c$$
$$\geq(k+m)\gamma-2\log(k+m)-c$$

for some constant c; clearly, this expression is greater than 0 from some k, m on. □

So, as a corollary of Propositions 1 and 2 we get the following statement.

Theorem 1. *For any code φ and any $\alpha\in(0,1]$ the Type I error of the test G_φ with level of significance α is not greater than α. If, in addition, the code φ is ideal then the Type II of G_φ error tends to 0 as the sample size n approaches infinity.*

3 Component independence testing

Component independence testing is the following task. A sample $Z=Z_1,\dots,Z_n$ is given where each Z_i consists of r components $Z_i^1, Z_i^2,\dots,Z_i^r$, $Z_i^j\in\mathbb{R}^{d_j}$. The sample is generated according to some probability distribution F_Z on $\mathbb{R}^d$, where $d:=\sum_{j=1}^r d_j$. The goal is to test whether the components are distributed independently. That is, H_0 is that

$$F_Z(Z_1^1\in T_1,\dots,Z_1^r\in T_r)=\prod_{j=1}^r F_Z(Z_1^j\in T_j) \tag{7}$$

for all measurable $T_j\subset\mathbb{R}^{d_j}$, $1\leq j\leq r$. H_1 is the negation of H_0 (the equality (7) is false for some selection of the sets T_j, $1\leq j\leq r$). Again, no assumption is made on the form of the distribution F_Z.

Fix any code φ and construct the test for component independence I_φ as follows. Assume that $n=2m$ for some m and define the samples X and $\bar{Y}$ as the first and the second half of the sample Z: $X_1=Z_1,\dots,X_m=Z_m$

and $\bar{Y}_1 = Z_{m+1}, \ldots, \bar{Y}_m = Z_{2m}$ (if n is odd then make samples X and $\bar{Y}$ of sizes $[n/2]$ and $n - [n/2]$). Construct the sample Y from $\bar{Y}$ by permuting the components independently: $Y_i^j = \bar{Y}_{\pi_j(i)}^j$, $1 \le i \le m$, $1 \le j \le r$ where π_j are permutations $1 \ldots m$, selected at random (with equal probabilities) independently of each other.

Definition 4 (Component independence test I_φ)**.** *The test I_φ (with level of significance α) consists in application of the test for homogeneity G_φ to the samples X and Y (with level of significance α).*

Indeed, it is easy to check that H_0 is true if and only if X and Y are distributed according to the same distribution. So we get the following statement.

Theorem 2. *For any code φ and any $\alpha \in (0, 1]$ the Type I error of the test I_φ with level of significance α is not greater than α. If, in addition, the code φ is ideal then the Type II error of I_φ error tends to 0 as the sample size n approaches infinity.*

References

[1] R. Cilibrasi, P. Vitányi, Clustering by Compression, IEEE Transactions on Information Theory 51(4) (2005).

[2] R. Cilibrasi, R. de Wolf , P. Vitányi, Algorithmic Clustering of Music, Computer Music Journal, 28(4) (2004) 49-67.

[3] E. Lehmann, Testing Statistical Hypotheses, 2nd edition, John Wiley & Sons, New York, 1986.

[4] M. Li, X. Chen, X. Li, B. Ma, P. Vitányi, The similarity metric, IEEE Trans. Inform. Th., 50(12) (2004), 3250- 3264.

[5] M. Li, P. Vitányi. An introduction to Kolmogorov complexity and its applications. Second edition, Springer, 1997.

[6] B. Ryabko, Prediction of random sequences and universal coding, Problems of Inform. Transmission, 24(2) (1988) 87-96.

[7] B. Ryabko, J. Astola, Universal Codes as a Basis for Time Series Testing, Statistical Methodology, 3, (2006) 375-397.

[8] B. Ryabko, V. Monarev, Using information theory approach to randomness testing, Journal of Statistical Planning and Inference, 133(1) (2005) 95-110.

[9] N. Vereshchagin, A. Shen and V. Uspensky. Lecture Notes on Kolmogorov Complexity, 2004, Unpublished, http://lpcs.math.msu.su/∼ver/kolm-book .

[10] J. Ziv, A. Lempel,Compression of individual sequences via variable-rate coding. IEEE Trans. Inform. Theory, vol. IT-24, no. 5, pp. 530-536, 1978.